arXiv:1605.00167v1 [cs.GT] 30 Apr 2016

Approximating Nash Equilibrium Via Multilinear Minimax

Bahman Kalantari
Department of Computer Science, Rutgers University, NJ
kalantari@cs.rutgers.edu

Abstract

On the one hand we state *Nash equilibrium* (NE) as a formal theorem on multilinear forms and give a pedagogically simple proof, free of game theory terminology. On the other hand, inspired by this formalism, we prove a *multilinear minimax theorem*, a generalization of von Neumann's bilinear minimax theorem. Next, we relate the two theorems by proving that the solution of a multilinear minimax problem, computable via linear programming, serves as an approximation to Nash equilibrium point, where its multilinear value provides an upper bound on a convex combination of *expected payoffs*. Furthermore, each positive probability vector once assigned to the set of players induces a *diagonally-scaled* multilinear minimax optimization with a corresponding approximation to NE. In summary, in this note we exhibit an infinity of multilinear minimax optimization problems each of which provides a polynomial-time computable approximation to Nash equilibrium point, known to be difficult to compute. The theoretical and practical qualities of these approximations are the subject of further investigations.

Keywords: Nash Equilibrium, von Neumann Minimax Theorem, Linear Programming, Matrix Scaling.

1 Nash Equilibrium

Nash equilibrium, [6], is a fundamental result in game theory. The complexity of its computing has been investigated extensively, e.g. in [1], [4], [5]. In this note we first formally state Nash equilibrium as a theorem on multidimensional matrices and give a proof not relying on game theory terminology (Section 1). On the one hand, pedagogically this simplifies its proof. On the other hand, it leads into a multidimensional generalization of von Neumann minimax theorem, where analogous to the bilinear case we give a proof via linear programming duality (Section 2). We then prove polynomial-time computable solution of a multidimensional minimax not only provides an approximation to Nash equilibrium point, but a nontrivial upper bound on a convex combination of expected payoffs in Nash equilibrium as well as the average expected payoffs (Section 3). Furthermore, we consider diagonal scaling of the multidimensional matrix to derive alternate bounds. In summary, solving multilinear minimax optimization problems is a way to compute approximate solutions to Nash equilibrium point in polynomial-time.

First, we give some notations and definitions. These are used throughout the article.

Notation 1. Given $n \in \mathbb{N}$, let $N = \{1, \ldots, n\}$. Given $n_i \in \mathbb{N}$, $i \in N$, let $N_i = \{1, \ldots, n_i\}$. Let $\widehat{N} = N_1 \times \cdots \times N_n$ and $\widehat{n} = n_1 \times \cdots \times n_n$. An n-dimensional $\widehat{n}$ matrix is a multidimensional array of real elements, written as $A = (a(I))$, where the index vector $I = (i_1, \ldots, i_n)$ ranges in $\widehat{N}$. For each $i \in N$, let $x^i = (x^i_1, \ldots, x^i_{n_i})^T \in \mathbb{R}^{n_i}$ be a vector of real variables. Set $x = (x^1, \ldots, x^n)$. Given $I = (i_1, \ldots, i_n) \in \widehat{N}$, we write $x(I)$ for the product $x^1_{i_1} \times \cdots \times x^n_{i_n}$. Consider the multilinear form corresponding to $A = (a(I))$:

$$A[x] \equiv A[x^1, \ldots, x^n] = \sum_{I \in \widehat{\mathbf{N}}} a(I)x(I).$$

For each $i \in N$, let $S_{n_i} = \{x \in \mathbb{R}^{n_i} : \quad \sum_{j=1}^{n_i} x_j = 1, \quad x \geq 0\}$. Let

$$\Delta_{\widehat{n}} = S_{n_1} \times \cdots \times S_{n_n}.$$

Given $p = (p^1, \ldots, p^n) \in \Delta_{\widehat{n}}$, $A[p]$ is the *value* of A at p. For each $x^i \in \mathbb{R}^{n_i}$, $i \in N$, define the linear form

$$A[p|x^i] = A[p^1, \ldots, p^{i-1}, x^i, p^{i+1}, \ldots, p^n].$$

Theorem 1. (Nash Equilibrium) *Let $A_0 = (a(i,I))$ be an $(n+1)$-dimensional $n \times \widehat{n}$ matrix, where $i \in N$ and $I \in \widehat{N}$. For each $i \in N$, consider the n-dimensional $\widehat{n}$ submatrix of A_0, $A_i = (a(i,I))$. Then, there exists $p_* = (p_*^1, \ldots, p_*^n) \in \Delta_{\widehat{n}}$ such that for each $i \in N$,*

$$\max_{x^i \in S_{n_i}} A_i[p_*|x^i] = A_i[p_*]. \tag{1}$$

Proof. For each $i \in N$, $j \in N_i$, let $e_j^{n_i} \in \mathbb{R}^{n_i}$ be the vector with 1 in its j-th coordinate and zero otherwise. From linearity, for each $i \in N$ we have

$$A_i[p_*|x^i] = \sum_{j=1}^{n_i} x_j^i A_i[p_*|e_j^{n_i}].$$

Thus, to prove (1) it is enough to prove there exists $p_* \in \Delta_{\widehat{n}}$ such that

$$A_i[p_*|e_j^{n_i}] \leq A_i[p_*], \quad \forall i \in N, \quad \forall j \in N_i. \tag{2}$$

Employing (2), an iterative scheme will be defined that given $p = (p^1, \ldots, p^n) \in \Delta_{\widehat{n}}$ that does not satisfy (2), attempts to find an improved point via the mapping $f(p) = (\overline{p}^1, \ldots, \overline{p}^n) \in \Delta_{\widehat{n}}$, defined as follows: For each $i \in N$ and $j \in N_i$, let

$$G_j^i(p) = \max\Big\{0,\ A_i[p|e_j^{n_i}] - A_i[p]\Big\}, \tag{3}$$

$$g_j^i(p) = p_j^i + G_j^i(p), \tag{4}$$

$$\overline{p}_j^i = \frac{g_j^i(p)}{\sum_{k \in N_i} g_k^i(p)}. \tag{5}$$

From (3)-(5) it follows that $f(p) \in \Delta_{\widehat{n}}$. Since f is continuous, by Brouwer fixed point theorem, there exists $p = (p^1, \ldots, p^n) \in \Delta_{\widehat{n}}$ satisfying $f(p) = p$. We show p satisfies

$$G_j^i(p) = 0, \quad \forall i \in N, \quad \forall j \in N_i. \tag{6}$$

Clearly (6) implies (2). From (4) and (5), and that $f(p) = p$, we have

$$p_j^i = \frac{p_j^i + G_j^i(p)}{\sum_{k \in N_i}(p_k^i + G_k^i(p))} = \frac{p_j^i + G_j^i(p)}{1 + \sum_{k \in N_i} G_k^i(p)}, \tag{7}$$

where the second equality uses that $\sum_{k \in N_i} p_k^i = 1$. Let $c_i = \sum_{k \in N_i} G_k^i(p)$. From (7) it follows that

$$c_i p_j^i = G_j^i(p), \quad \forall j \in N_i. \tag{8}$$

Suppose p does not satisfy (1). Then there exists $i \in N$ such that $c_i > 0$. From (3) and (8) we have

$$c_i p_j^i = A_i[p|e_j^{n_i}] - A_i[p], \quad \forall j \in N_i. \tag{9}$$

Multiplying both sides of (9) by p_j^i, summing over $j \in N_i$ and using linearity of $A_i[p|e_j^{n_i}]$, we get

$$c_i \sum_{j \in N_i} (p_j^i)^2 = \sum_{j=1}^{n_i} p_j^i A_i[p|e_j^{n_i}] - \sum_{j=1}^{n_i} p_j^i A_i[p] = A_i[p] - A_i[p] = 0. \tag{10}$$

But the left-hand-side of (10) is positive, a contradiction. Hence the proof of (6). □

1.1 Game Theory Terminology

Using the same notation as above, N is the set of *players*. For each player $i \in N$, N_i is his set of *pure actions* or *pure strategies*. Each $I = (i_1, \ldots, i_n) \in \widehat{N}$ is a *pure strategy profile*. For each $I \in \widehat{N}$, the matrix entry $a(i, I)$ is *payoff* or *utility* for player i when pure strategy profile I is selected. Each player i may choose an action according to a *mixed strategy* with probability vector $p^i \in S_{n_i}$. When a profile is selected with probability vector $p = (p^1, \ldots, p^1_n) \in \Delta_{\widehat{n}}$, the *expected utility* of player i is $A_i[p]$. Nash equilibrium states that there exists an *equilibrium point* $p_* = (p^1_*, \ldots, p^n_*) \in \Delta_{\widehat{n}}$ such that if player k knows that each players $i \neq k$ will use probability vectors p^i_*, he cannot improve his expected utility by selecting any other probability vector in S_{n_k} than p^k_*. There is thus equilibrium.

The question of complexity of computing or approximating the equilibrium point $p_* \in \Delta_{\widehat{n}}$ has been studied extensively, [1], [4], [5]. Even its approximation is hard. One may ask: Can we compute upper bounds on $A_i[p_*]$? There are of course trivial upper bounds such as the maximum entry of the the matrix A_0, however we may ask: Can we compute nontrivial upper bounds on $A_i[p_*]$? Motivated by this question, in the next section we prove a multilinear form of von Neumann's minimax theorem which is an interesting problem of independent interest. Then in Section 3 we show how an optimal solution of the minimax problem gives a bound on a convex combination of $A_i[p_*]$'s, their average, as well as an approximation to p_*. In Section 4, we consider some ideas from diagonal scaling to derive alternate bounds on a convex combination of $A_i[p_*]$, as well as different approximations to p_*.

2 A Mulitinear von Neumann Minimax Theroem

In this section we state and prove a multilinear minimax theorem. As before, let $A_0 = (a(i, I))$, $i \in N$, $I = (i_1, \ldots, i_n) \in \widehat{N}$ be the $(n+1)$-dimensional matrix. Given $x \in S_n$ and $p \in \Delta_{\widehat{n}}$, consider the multilinear from

$$A_0[x, p] = \sum_{i \in N} \sum_{I \in \widehat{N}} x_i a(i, I) p(I), \tag{11}$$

where $p(I) = p^1_{i_1} \times \cdots \times p^n_{i_n}$. Clearly,

$$A_0[x, p] = x_1 A_1[p] + \cdots + x_n A_n[p]. \tag{12}$$

Theorem 2. (Multilinear von Neumann Minimax Theorem)

$$\min_{x \in S_n} \max_{p \in \Delta_{\widehat{n}}} A_0[x, p] = \max_{p \in \Delta_{\widehat{n}}} \min_{x \in S_n} A_0[x, p]. \tag{13}$$

Proof. The proof of (bilinear) von Neumann minimax can be established via linear programming duality theory, see e.g. [2], [3]. Analogously, we will represent the left-hand-side and right-hand-side of (13) as a pair of primal and dual linear programs. However, the proof is somewhat trickier than the bilinear case.

For each $x \in S_n$, the maximum of $A_0[x, p]$ over $p \in \Delta_{\widehat{n}}$ is the maximum of $\sum_{i \in N} a(i, I) x_i$, as I ranges in $\widehat{N}$. Thus the value of the left-hand-side in (13) is equivalent to the optimal value of the following *primal* LP:

$$(LP) \qquad \begin{aligned} &\min \delta \\ &\sum_{i \in N} a(i, I) x_i \leq \delta, \quad \forall I \in \widehat{N}, \\ &\sum_{i \in N} x_i = 1, \\ &x_i \geq 0, \quad \forall i \in N. \end{aligned}$$

Next, we consider the value of the right-hand-side in (13). For each $p \in \Delta_{\widehat{n}}$, the minimum of $A_0[x, p]$ over $x \in S_n$ is the minimum of $\sum_{I \in \widehat{N}} a(i, I) p(I)$ as i ranges in N. We can find the maximum of these over

$p \in \Delta_{\widehat{n}}$ by putting a lower bound of λ on each such term and then maximize it over all I in $\widehat{N}$. For each $I \in \widehat{N}$ we introduce a variable $q(I)$ corresponding to the term $p(I)$. Now consider the following *dual* LP.

$$
\begin{aligned}
(DLP) \qquad & \max \lambda \\
& \sum_{I \in \widehat{N}} a(i,I) q(I) \geq \lambda, \quad \forall i \in N, \\
& \sum_{I \in \widehat{N}} q(I) = 1, \\
& q(I) \geq 0, \quad \forall I \in \widehat{N}.
\end{aligned}
$$

We prove (DLP) is the dual of (LP). Let A be the $\widehat{n} \times n$ constraint matrix corresponding to the inequalities in (LP) (excluding $x \geq 0$). Let $e_n \in \mathbb{R}^n$, and $e_{\widehat{n}} \in \mathbb{R}^{\widehat{n}}$ be the vector of ones. Then (LP) is equivalent to

$$\min\{\delta : \ Ax \leq \delta e_{\widehat{n}}, \ e_n^T x = 1, \ x \geq 0\}. \tag{14}$$

Introducing slacks, s, and writing δ as the difference of two nonnegative variables, (14) is equivalent to:

$$\min\{\delta_1 - \delta_2 : \ Ax - \delta_1 e_{\widehat{n}} + \delta_2 e_{\widehat{n}} + s = 0, \ e_n^T x = 1, \ x \geq 0\}. \tag{15}$$

Since (15) is an stand form LP, its dual can easily be seen to be

$$\max\{\lambda : \ A^T w + \lambda e_n \leq 0, \ -e_{\widehat{n}}^T w \leq 1, \ e_{\widehat{n}}^T w \leq -1, \ w \leq 0\}. \tag{16}$$

Replacing w with $-w$, (16) reduces to

$$\max\{\lambda : \ A^T w \geq \lambda e_n, \ e_{\widehat{n}}^T w = 1, \ w \geq 0\}. \tag{17}$$

By identifying w in (17) with the vector of variables corresponding to $q(I)$, $I \in \widehat{N}$, it follows that (17) is equivalent to (DLP). Hence (LP) and (DLP) are primal-dual pair.

Let us consider the set of $q(I)$, $I \in \widehat{N}$, satisfying the constraints $\sum_{I \in \widehat{N}} q(I) = 1$ as an ordered vector, $Q(\widehat{N})$ in $S_{\widehat{n}}$, the unit simplex in dimension $\widehat{n}$. Next, for each $Q(\widehat{N}) \in S_{\widehat{n}}$ we construct a unique $p \in \Delta_{\widehat{n}}$, to be referred as its *derived point*: For each $i \in N$ and each $j \in N_i$, let $\widehat{N}_j^i$ be the set of all $I \in \widehat{N}$ whose i-th coordinate is j. Then set

$$p_j^i = \sum_{I \in \widehat{N}_j^i} q(I). \tag{18}$$

It is easy to check that the vector $p^i = (p_1^i, \ldots, p_{n_i}^i)^T$ lies in S_{n_i}. Thus setting $p = (p^1, \ldots, p^n)$, it follows that $p \in \Delta_{\widehat{n}}$. Hence the proof of theorem. □

Corollary 1. *Let x^* be an optimal solution of (LP). Let $Q^*(\widehat{N}) = \{q^*(I) : I \in \widehat{N}\} \in S_{\widehat{n}}$ be an optimal solution of (DLP) and let $p^* \in \Delta_{\widehat{n}}$ be the corresponding derived solution (see (18)). Then*

$$\min_{x \in S_n} \max_{p \in \Delta_{\widehat{n}}} A_0[x,p] = \max_{p \in \Delta_{\widehat{n}}} \min_{x \in S_n} A_0[x,p] = A_0[x^*, p^*]. \tag{19}$$

Furthermore, for each $x \in S_n$ and each $p \in \Delta_{\widehat{n}}$ we have

$$A_0[x^*, p] \leq A_0[x^*, p^*] \leq A_0[x, p^*]. \tag{20}$$

Proof. By the fact that (LP) and (DLP) are primal and dual pair, it follows that if δ^* and λ^* are the respective optimal objective values, then $\delta^* = \lambda^*$.

Given $Q(\widehat{N}) \in S_{\widehat{n}}$, let $p \in \Delta_{\widehat{n}}$ be its derived point. Next substitute x^* for x and δ^* for δ in (LP). Then for each $I \in \widehat{N}$ multiply the corresponding inequality constraint by $q(I)$, then sum over these inequalities while replacing $Q(\widehat{N})$ by its derived point to get,

$$A_0[x^*, p] \leq \delta^* \sum_{I \in \widehat{N}} q(I) = \delta^*. \tag{21}$$

In particular, substituting $Q^*(\widehat{N})$ we get

$$A_0[x^*, p^*] \leq \delta^*. \tag{22}$$

Next substitute for $q^*(I)$, for $q(I)$, $I \in \widehat{N}$ in the corresponding inequality in (DLP), then for $i \in N$, multiply the i-th inequality constraint by x_i, replace $Q^*(\widehat{N})$ by its derived point p^* and finally sum these up to get

$$A_0[x, p^*] \geq \lambda^* \sum_{i \in N} x_i = \lambda^*. \tag{23}$$

In particular, substituting x^* we get

$$A_0[x^*, p^*] \geq \lambda^*. \tag{24}$$

Now from (22) and (24) and that $\delta^* = \lambda^*$ we get (19). Next from (21), (23) and (19) we get (20). □

3 Bounds On Expected Payoffs in Nash Equilibrium

We now state a connection between Nash equilibrium and an optimal von Neumann minimax pair.

Theorem 3. *Let p_* be an optimal Nash equilibrium point and (x^*, p^*) an optimal pair of multilinear von Neumann minimax. We have*

$$A_0[x^*, p_*] \leq A_0[x^*, p^*] = \min_{i \in N} A_i[p^*]. \tag{25}$$

Let

$$\sigma_n = \frac{1}{n} \sum_{i=1}^{n} A_i[p_*].$$

Let $x^*_{\min} = \min\{x^*_i : i = 1, \ldots, n\}$. If $x^*_{\min} > 0$, and $A_i[p_*] \geq 0$ for all $i = 1, \ldots, n$, then

$$\sigma_n \leq \frac{1}{n x^*_{min}} A_0[x^*, p^*]. \tag{26}$$

Proof. The inequality in (25) follows from (20) in Corollary 1, selecting $p = p_*$. To show the equality in (25), from (20) and selecting $x = e^n_i$, the n-vector with one in its i-th position and zero otherwise, it follows that $A_0[x^*, p^*] \leq \min\{A_i[p^*] : i = 1, \ldots, n\}$. On the other hand,

$$A_0[x^*, p^*] = \sum_{i=1}^{n} x^*_i A_i[p^*] \geq \sum_{i=1}^{n} \min\{A_i[p^*] : i = 1, \ldots, n\}.$$

To prove (26), we have

$$A_0[x^*, p_*] = \sum_{i=1}^{n} x^*_i A_i[p_*] \geq n x^*_{min} \sigma_n.$$

Hence the proof. □

Theorem 3 implies that by computing the von Neumann minimax pair (x^*, p^*) we can derive an upper bound on a convex combination of expected payoffs of players at the Nash equilibrium point, as well as the average expected payoffs. We can thus consider p^* as an approximation to p_*. The quality of any approximation to p_* can be quantified is as follows.

Definition 1. We say $p \in \Delta_{\widehat{n}}$ is a t-approximate Nash equilibrium, if $t = \max\{t_i : i = 1, \ldots, n\}$, where for each $i = 1, \ldots, n$, t_i is defined as

$$\max_{x^i \in S_{N_i}} A_i[p|x^i] = t_i A_i[p].$$

Clearly, the closer t is to one, the better p approximates p_*, or more precisely, $A_i[p]$ approximates $A_i[p_*]$. The question arises, how good is the quality of p^*? This, we feel, is an important question open to further investigation.

In the next section we consider alternate bounds by solving alternate minimax problems.

4 Alternate Bounds Via Diagonally Scaled Minimax Relaxations

In what follows we consider alternate minimax problems derived via diagonal scaling of the matrix $A_0 = (a(i, I))$. Let

$$S_n^\circ = \{x \in S_n : x > 0\}.$$

Given each $d = (d_1, \ldots, d_n) \in S_n^\circ$, set

$$A_0(d) = (d_i a(i, I)).$$

This corresponds to scaling each submatrix A_i of A_0 by d_i. Thus $A_i(d) = d_i A_i$. We can view d as a vector of positive probabilities assigned to the set of n players. Next consider the optimal *diagonally-scaled* minimax pair of solutions for the matrix $A_0(d)$, denote the pair by $(x^*(d), p^*(d))$. In particular, (x^*, p^*) in the previous section corresponds to $d = (1/n, \ldots, 1/n)^T \in \mathbb{R}^n$. As before, denoting p_* as the Nash equilibrium point for A_0, note that it also remains an equilibrium point for $A_0(d)$. However, from Theorem 3 we have

$$A_0(d)[x^*(d), p_*] \leq A_0(d)[x^*(d), p^*(d)].$$

Equivalently,

$$\sum_{i \in N} d_i x_i^*(d) A_i[p_*] \leq \sum_{i \in N} d_i x_i^*(d) A_i[p^*(d)].$$

Dividing both sides by $\sigma = \sum_{i \in N} d_i x_i^*(d)$ and letting $d_i' = d_i x_i^*(d)/\sigma$, from $d \in S_n^\circ$ we get a new point $d' \in S_n$ satisfying an alternate bound

$$\sum_{i \in N} d_i' A_i[p_*] \leq \sum_{i \in N} d_i' A_i[p^*(d)].$$

If $d' > 0$, we can replace d with d' and repeat the process. Replacing d with d' corresponds to a projective transformation. Some questions arise regarding the above bound for arbitrary $d \in S_n^\circ$. What if we randomly select such d and compute the above? If $d' = e_i^n$, then we will have a bound on $A_i[p_*]$. In other words we may be able to produce bounds on individual player payoffs. Another question is what will be the infimum of the above over $d \in S_n^\circ$? From the practical point of view, for each $d \in S_n^\circ$ the computation of $p^*(d)$ gives rise to a $t(d)$-approximate solution (see Definition 1). Ideally, the infimum of $t(d)$ equal one.

Final Remarks

Our initial goal was to state a formal but pedagogically simple proof of Nash equilibrium, free of its game theory interpretations. Numerous proofs are available on the Internet. However, most proofs seem to obscure the mathematical properties behind the game theory terminology. Nash's theorem, despite its tremendous ingenuity and applications, after all is a theorem on multilinear forms. This formal view has offered new insights on Nash equilibrium. In particular, by proving a multilinear minimax theorem we have arrived at polynomial-time approximations to Nash equilibrium. These are worthy of further investigation, especially in view of the inherent difficulty in computing NE, as revealed by tremendous research in game theory.

As shown here, by solving a primal-dual pair of linear programs, it is possible to obtain upper bounds on a weighted expected payoff in Nash equilibrium, the average expected payoff, as well as approximate solutions to equilibrium point. Furthermore, diagonal scaling of the multidimensional matrix gives rise to new upper bounds and approximate solutions. The fact that the multilinear minimax problem, considered as a relaxations of Nash equilibrium, is solvable in polynomial-time makes it tractable in theory. It is also an interesting problem in its own right. For instance, it is interesting to note that since (DLP) has $(n+1)$ nontrivial constraints, in an optimal solution at most $(n+1)$ of the $q(I)$'s are positive. Note that even solving the multilinear minimax as a relaxation of Nash equilibrium could be a difficult problem when there are several players, each having $O(n)$ strategies. This is because in such case the bound based on the complexity of polynomial-time LP algorithms is a high degree polynomial in n. Nevertheless, having obtained the minimax pair of solutions (x^*, p^*), we may consider p^* as a relaxation to p_* and measure its quality as a t-approximation solution (see Definition 1). In this note we have also discussed ways to obtain different approximate solutions through diagonal scaling, resulting in different t-approximate solutions.

Several questions arise, such as: (1) How efficiently can the minimax problems be solved? That is, is there a way to make use of their structure to give faster algorithms than the straightforward application of known LP algorithms? (2) Can we bound the quality of p^* as a t-approximate solution to p_*? (3) Are there classes of multidimensional matrices for which p^* gives a good t-approximate solution to p_*? (4) What is the quality of approximation under diagonal scalings? (5) Finally, can we derive lower bounds to Nash equilibrium expected payoffs? e.g. instead of solving a minimax problem, by solving a maximin problem over the corresponding domains.

References

[1] X. Chen and X. Deng, Settling the complexity of 2-player Nash equilibrium, *In the 47-th Annual IEEE Symposium on Foundations of Computer Science*, FOCS, 2006.

[2] V. Chvátal, *Linear Programming*, W.H. Freeman and Company, New York, 1983.

[3] G. B. Dantzig. *Linear Programming and Extensions*, Princeton University Press, 1963.

[4] C. Daskalakis, P.W. Goldberg and C.H. Papadimitriou, The complexity of computing Nash equilibrium, *In the 38-th Annual Symposium on Theory of Computing*, STOC, 2006.

[5] K. Etessami and M. Yannakakis, On the complexity of computing Nash equilibria and other fixed points, *In the 48-th Annual IEEE Symposium on Foundations of Computer Science*, FOCS, 2007.

[6] J. Nash, Noncooperative games, *Annals of Mathematics*, 54, 289-295, 1951.

[7] J. von Neumann, Zur theorie der gesellshaftsspiele, *Mathematische Annalen*, 100: 295320, 1928.